OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

GronOR application readiness for Summit

2020 Performance, Portability, and Productivity in HPC Forum
September 2, 2020

Tjerk P. Straatsma
Distinguished Research Scientist, ORNL
Adjunct Professor of Chemistry, UA

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

Application Readiness on Summit

Application readiness criteria:

1. Scalability to at least 20% of the number of Summit nodes

2. Performance improvement with factor 2 or better from utilizing the GPUs.

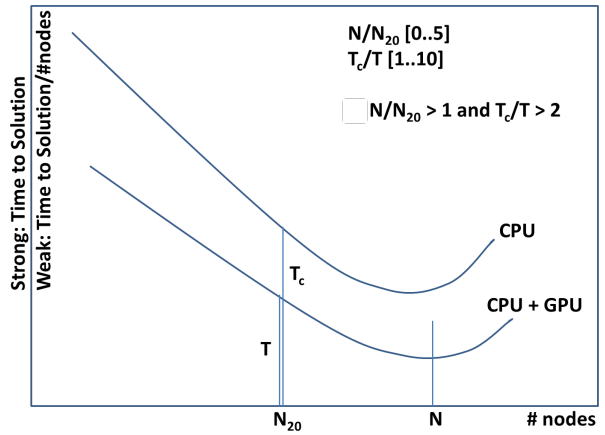

Application	Scaling	Acceleration @20%	Program	Petascale
LS-Dalton	1024	2.1	CAAR	
FLASH	4592	3.1	CAAR	
RAPTOR	3888	4.0	CAAR	Titan
NUCCOR	1024	4.1	CAAR	
SpecFEM	960	4.9	CAAR	
NWChem	2048	5.0	CAAR	Titan
NAMD	1024	5.7	CAAR	Titan
E3SM	1024	8.1	CAAR	Titan
XGC	2018	13.2	CAAR	Titan
GronOR	4604	14.8	ESP	
GTC	928	20.0	CAAR	Titan
QMCPack	4096	24.7	CAAR	Titan
HACC	1024	30	CAAR	Titan
DIRAC	1024	46	CAAR	

GronOR Non-Orthogonal Configuration Interaction Methodology

- Wave functions Ψ are expanded in terms of many-electron basis functions Φ_i that can be Slater determinants or linear combinations thereof:

$$\Psi = \sum_{i=1}^{N} c_i \Phi_i$$

- Slater determinants Φ_i are anti-symmetrized products of orbitals φ_j:

$$\Phi_i = \mathrm{A} \prod_j \varphi_j$$

- Orbitals φ_j, and consequently the Slater determinants Φ_i , do not need to be orthogonal:

$$\langle \Phi_i | \Phi_j \rangle = S_{ij}$$

- This non-orthogonality complicates calculation of required Hamiltonian matrix elements

$$\langle \Phi_i | H | \Phi_j \rangle = H_{ij}$$

- Solve (**H**-*E***S**)(**c**)=0 to get energies and Ψ (i.e. the expansion coefficients c_i)

Advantages of Non-Orthogonal Configuration Interaction

- Wave function expansions allow for the inclusion of:
  - Orbital relaxation effect
  - Correlation effects

- Expansion in terms of non-orthogonal determinants Φ_i lead to shorter expansions

- Wave functions in terms of non orthogonal orbitals:
  - Easy description of systems in terms of individually optimized states of system
  - Easy description of systems in terms of predefined states of system components
  - Facilitates chemical interpretation

GNOME Factorization of Cofactors

- Two electron matrix elements between non-orthogonal determinants Δ_a and Δ_b:

$$\mathrm{I}_2 = \sum_{k>i}^{N} \sum_{l>j}^{N} < aiak|O_{12}|b_k b_l > S(ik,jl)$$

- Expansion in atomic basis functions {χ} and {χ'}, and factorization of the cofactor matrix leads to:

$$\mathrm{I}_2 = \sum_{r>p}^{m} \sum_{s>q}^{n} < \chi_p \chi_r|O_{12}|\chi'_q \chi'_s > B(pr,qs)$$

$$= \sum_{r>p}^{m} \sum_{s>q}^{n} < \chi_p \chi_r|O_{12}|\chi'_q \chi'_s > (1-ppr)(1-pqs)F(\omega)_{pq} G(\omega)_{rs}$$

- The number of singularities ω determines the functional form of F and G in terms of the expansion coefficients

Non-Orthogonal Configuration Interaction

- Description of molecular assemblies in terms of combinations of molecular electronic states
  - Resembles valence bond picture of electronic structure in terms of Lewis structures
  - Intuitive description of local processes, such as excitations of one molecule

- Put in practice using the embedded cluster model:
  1. Generation of many electron basis functions (MEBFs) as spin-adapted anti-symmetrized combinations (SAACs) of molecular wave functions
  2. Non-Orthogonal Configuration Interaction (NOCI) calculation

- Wave functions can be of any kind, e.g. correlated wave functions
  - Complete Active Space Self Consistent Field (CASSCF), or
  - Configuration Interaction (CI) calculations
  - Fully optimized molecular state wave functions lead to:
    - Non-orthogonal MEBFs
    - Proper inclusion of orbital relaxation and local correlation effects
    - Allows direct calculation of interaction matrix elements between different electronic states
    - Proper description possible for physical processes such as
      - Charge separation following photo-excitation
      - Singlet fission of photo-excited state on one molecule and two molecular triplet states

Example Application: Singlet Fission

A potential application for GronOR is to describe singlet fission

In singlet fission, the electronic coupling between S_0S_1 and ^1TT states is quantity of interest.

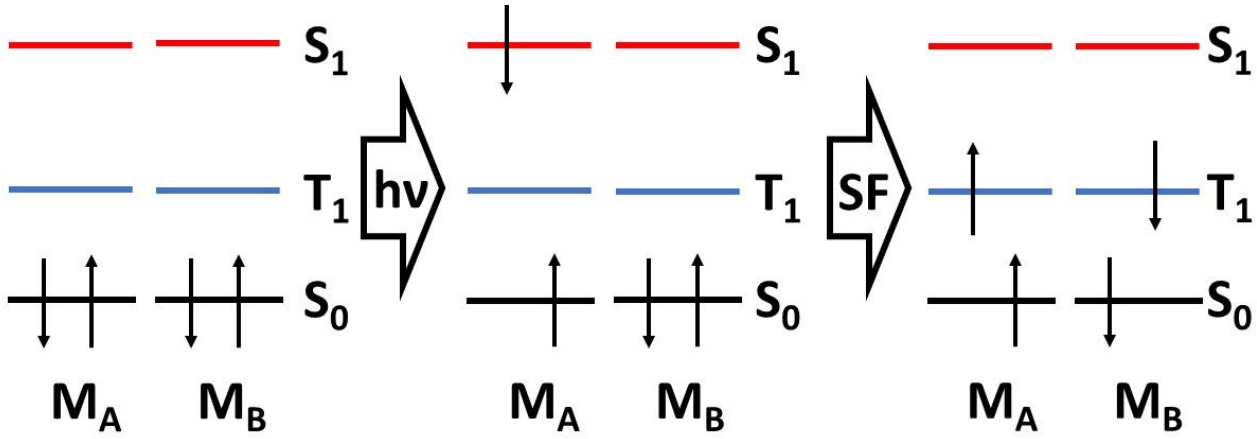

Spin-adapted product wave functions

- Consider two molecules for which the S_0, S_1, T_1, D^+_0 and D^-_0 molecular states have been calculated
- Leads to possible formation of seven spin singlet combinations
- MEBFs built from molecular wave functions as linear combinations of anti-symmetrized products of the different molecular M_S functions, such that the total M_S equals 0.

$$\left|\Phi_1\right\rangle = \left|\Psi_A^{S_0}\Psi_B^{S_0}\right\rangle \qquad \left|\Phi_5\right\rangle = \left|\Psi_A^{T_1}\Psi_B^{T_1}\right\rangle$$

$$\left|\Phi_2\right\rangle = \left|\Psi_A^{S_0}\Psi_B^{S_1}\right\rangle \qquad \left|\Phi_6\right\rangle = \left|\Psi_A^{D_0^+}\Psi_B^{D_0^-}\right\rangle$$

$$\left|\Phi_3\right\rangle = \left|\Psi_A^{S_1}\Psi_B^{S_0}\right\rangle \qquad \left|\Phi_7\right\rangle = \left|\Psi_A^{D_0^-}\Psi_B^{D_0^+}\right\rangle$$

$$\left|\Phi_4\right\rangle = \left|\Psi_A^{S_1}\Psi_B^{S_1}\right\rangle$$

- For example, molecular triplet states are combined into a singlet according to

$$\left|\Phi_5\right\rangle = -\frac{1}{3}\sqrt{3}\hat{A}\left|\left(\Psi_A^T\right)_{M_S=+1}\left(\Psi_B^T\right)_{M_S=-1}\right\rangle - \frac{1}{3}\sqrt{3}\hat{A}\left|\left(\Psi_A^T\right)_{M_S=-1}\left(\Psi_B^T\right)_{M_S=+1}\right\rangle + \frac{1}{3}\sqrt{3}\hat{A}\left|\left(\Psi_A^T\right)_{M_S=0}\left(\Psi_B^T\right)_{M_S=0}\right\rangle$$

- The MEBF built from two molecular doublet states is a linear combination of anti-symmetrized products of the M_S = +½ and M_S =-½, and M_S = -½ and M_S = +½ molecular functions

$$\left|\Phi_6\right\rangle = \frac{1}{2}\sqrt{2}\hat{A}\left|\left(\Psi_A^{D_0^+}\right)_{M_S=+1/2}\left(\Psi_B^{D_0^-}\right)_{M_S=-1/2}\right\rangle - \frac{1}{2}\sqrt{2}\hat{A}\left|\left(\Psi_A^{D_0^+}\right)_{M_S=-1/2}\left(\Psi_B^{D_0^-}\right)_{M_S=+1/2}\right\rangle$$

Spin-adapted product wave functions

- The final NOCI wave function is then written as a linear combination of the MEBF

$$\Psi_{NOCI} = \sum_{\mu} C_{\mu} \left| \Phi_{\mu} \right\rangle$$

- Coefficients C_{μ} are determined in the usual way with variational theory
- Determination of C_{μ} requires the Hamiltonian matrix elements and overlap matrix elements

$$\left\langle \Phi_{\mu} \right| H \left| \Phi_{\nu} \right\rangle \qquad \left\langle \Phi_{\mu} \middle| \Phi_{\nu} \right\rangle$$

- Use of optimized orbitals for each molecular electronic state leads to non-orthogonal orbitals in

$$\left| \Psi_A^I \right\rangle_{\sigma} \quad \text{and} \quad \left| \Psi_A^J \right\rangle_{\sigma'}$$

GronOR Computational Complexity

- For an assembly consisting of two molecules, where the molecular CASSCF wavefunctions consist of 500 determinants:

- The MEBFs have ~10^5 determinants
- The number of matrix elements over non-orthogonal determinants to be calculated is ~10^{10}

- Many of these determinants will be zero

- If the product of the coefficients $c_{\mu\sigma}^{IAJB} c_{\nu\sigma}^{KALB}$ is small, the matrix elements do not need to be evaluated

- In the example this would reduce the number of matrix elements to ~10^7

GronOR Implementation

- Currently interfaced with:
  - OpenMolcas to obtain the CASSCF CI vector and the state specific CASSCF orbitals
  - OpenMolcas to obtain the required two-electron integrals

- First step is the generation of the anti-symmetrized product determinants with products of coefficients
  - Using appropriate spin-coupling coefficients, the products are combined into anti-symmetrized spin-adapted combinations
  - This step is not parallelized, as it takes minimal time

- Second step is the evaluation of Hamiltonian and overlap matrix elements
  - Requires the processing of a large set of two-electron integrals
  - Is the computationally most time consuming part of the calculation
  - Involves many contributions in the form of determinant pairs that can be calculated independently

GronOR Parallel Accelerated Implementation

- Massively parallel implementation of the algorithm adopts a task-based approach
  - Scheduler/Execution processing model with process groups evaluating batches of matrix elements
  - Current implementation allows batched execution of numbers of matrix elements
  - Work between execution process groups is naturally load-balanced
  - Execution process group size determined by memory needed to hold 2-electron integrals
  - Computationally extensive work is GPU accelerated, using OpenACC directives and CUSOLVER library
  - Asynchronous use of host and accelerator in evaluating matrix elements
  - Implementation is hard-fault resilient

GronOR Parallel Implementation

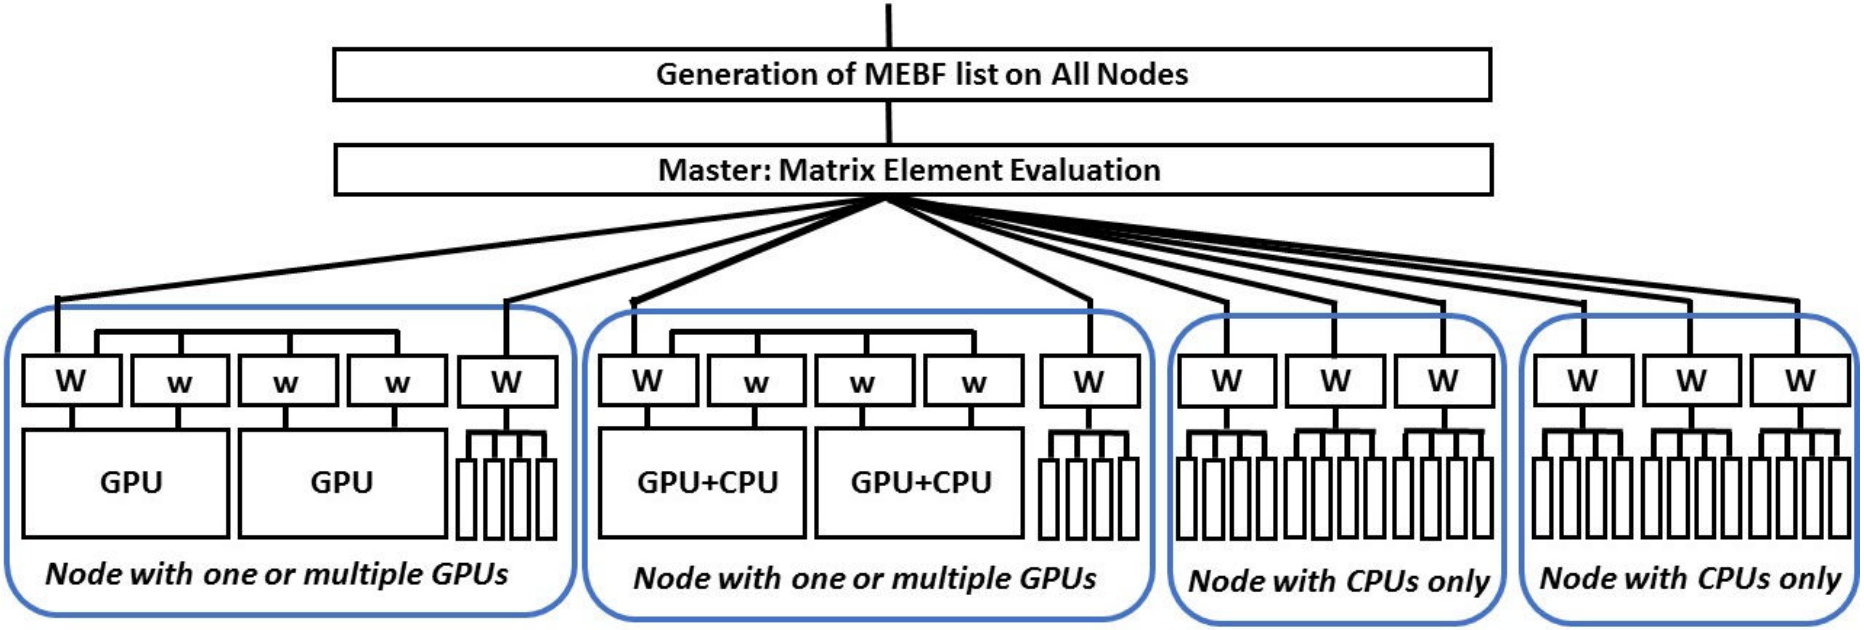

W: worker communicating with master, w: worker communicating with first worker in a group

Implementation details:

- Different node architectures can be used in a single calculation
- GPU accelerators can be shared between worker processes
- Worker processes can use CPU (single thread) and GPU asynchronously in single matrix element calculation
- Worker processes without GPU access can use OpenMP threading on multiple cores

GronOR Scheduler-Execution Communication

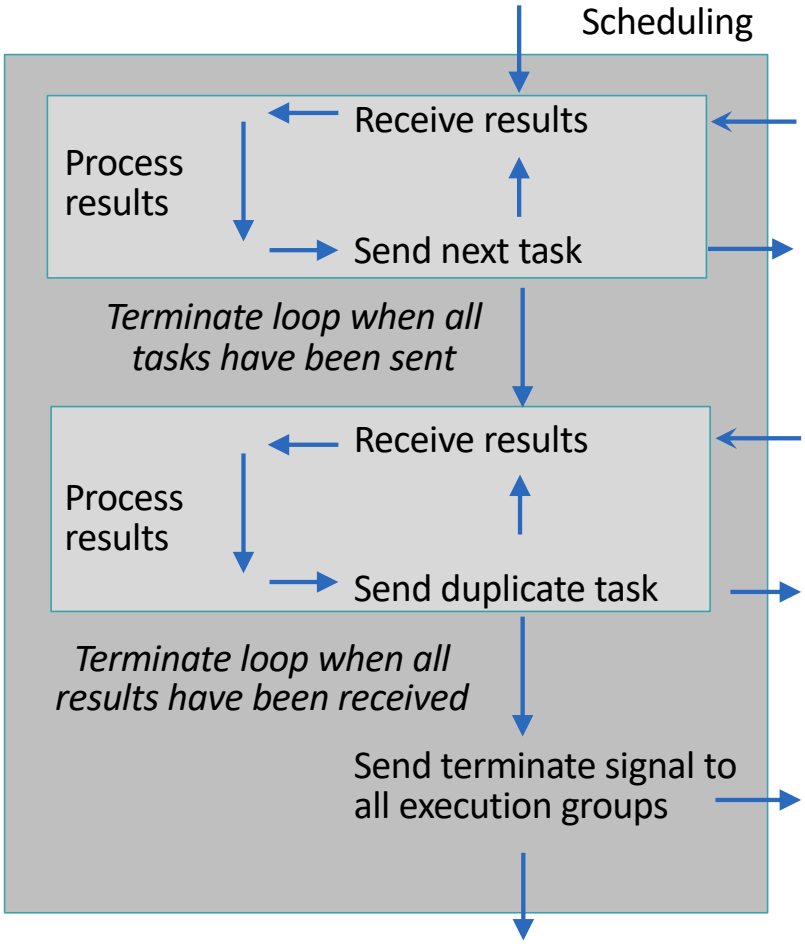

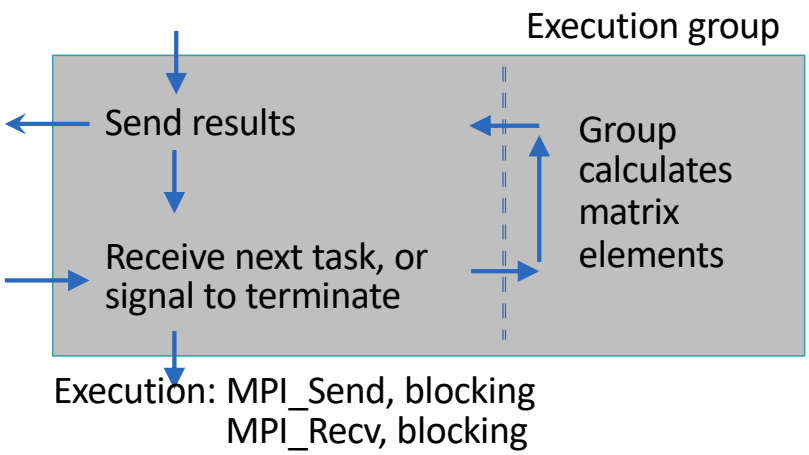

Execution: MPI_Send, blocking
MPI_Recv, blocking

Scheduler: MPI_Recv, blocking from ANY_SOURCE
MPI_Send, blocking

MPI_iSend, non-blocking to terminate

GronOR Benchmark: Two Naphthalene Molecules

Naphthalene itself is not interesting for singlet fission, a dimer of naphthalene molecules can serve perfectly well for doing performance tests

Geometry was optimized at the DFT level (B3LYP)

Considered four MEBFs: $\left|\Psi_A^{S_0}\Psi_B^{S_0}\right\rangle \left|\Psi_A^{S_1}\Psi_B^{S_0}\right\rangle \left|\Psi_A^{S_0}\Psi_B^{S_1}\right\rangle \left|\Psi_A^{T_1}\Psi_B^{T_1}\right\rangle$

CASSCF calculations on the molecular S_0, S_1 ($^1B_{1u}$), and T_1 ($^3B_{1u}$) states

Active space consisting of:
four electrons in four orbitals (CAS(4,4)), and
eight electrons in eight orbitals (CAS(8,8)).

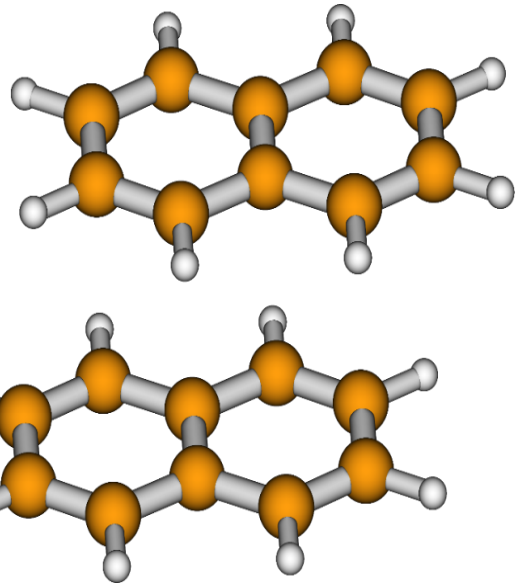

GronOR Benchmark: Two Naphthalene Molecules CAS(8,8)/CAS(4,4)

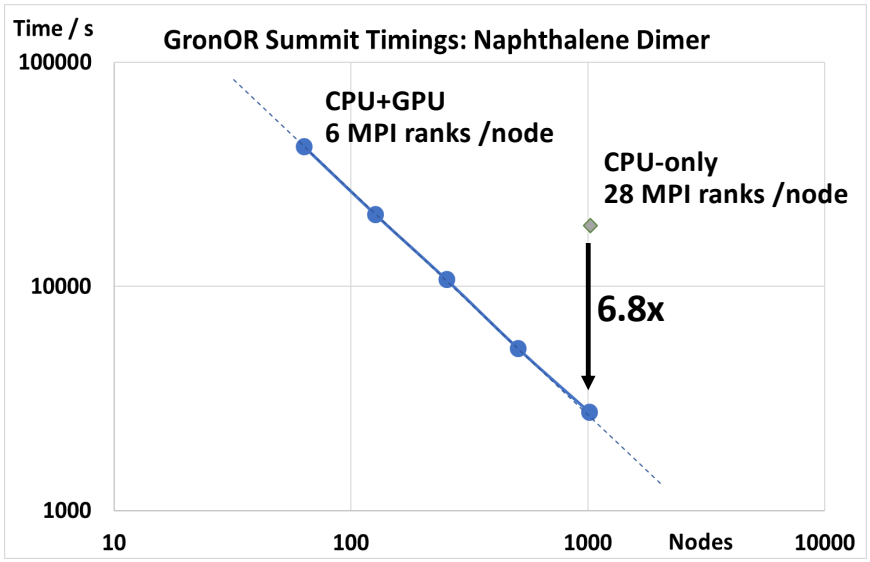

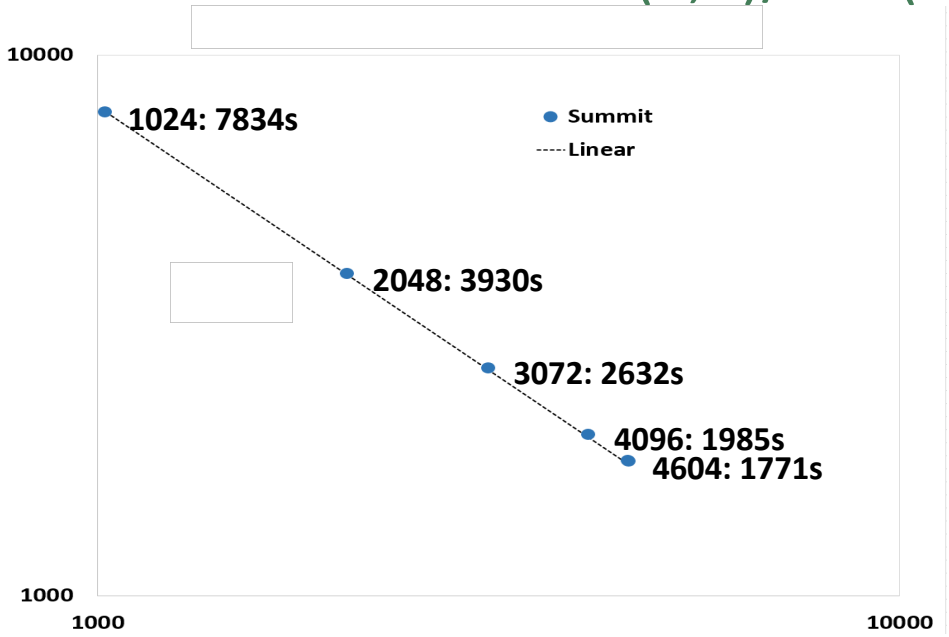

Instead of full 4x4 Hamiltonian a single matrix element $\left\langle \Psi_A^{S_0}\Psi_B^{S_0} \middle| H \middle| \Psi_A^{S_0}\Psi_B^{S_0} \right\rangle$ is calculated

This single element requires evaluating 112,867,800 matrix elements!!

The full 4x4 Hamiltonian requires evaluating 2,135,997,480 matrix elements!!!

Near-linear accelerated strong scaling benchmarks on Summit for 64 to 1024 nodes

GPU accelerated speedup is a factor of 6.8 comparing six MPI+GPU and 28 MPI+CPU-only

GronOR Power Trace on 4590 Summit nodes

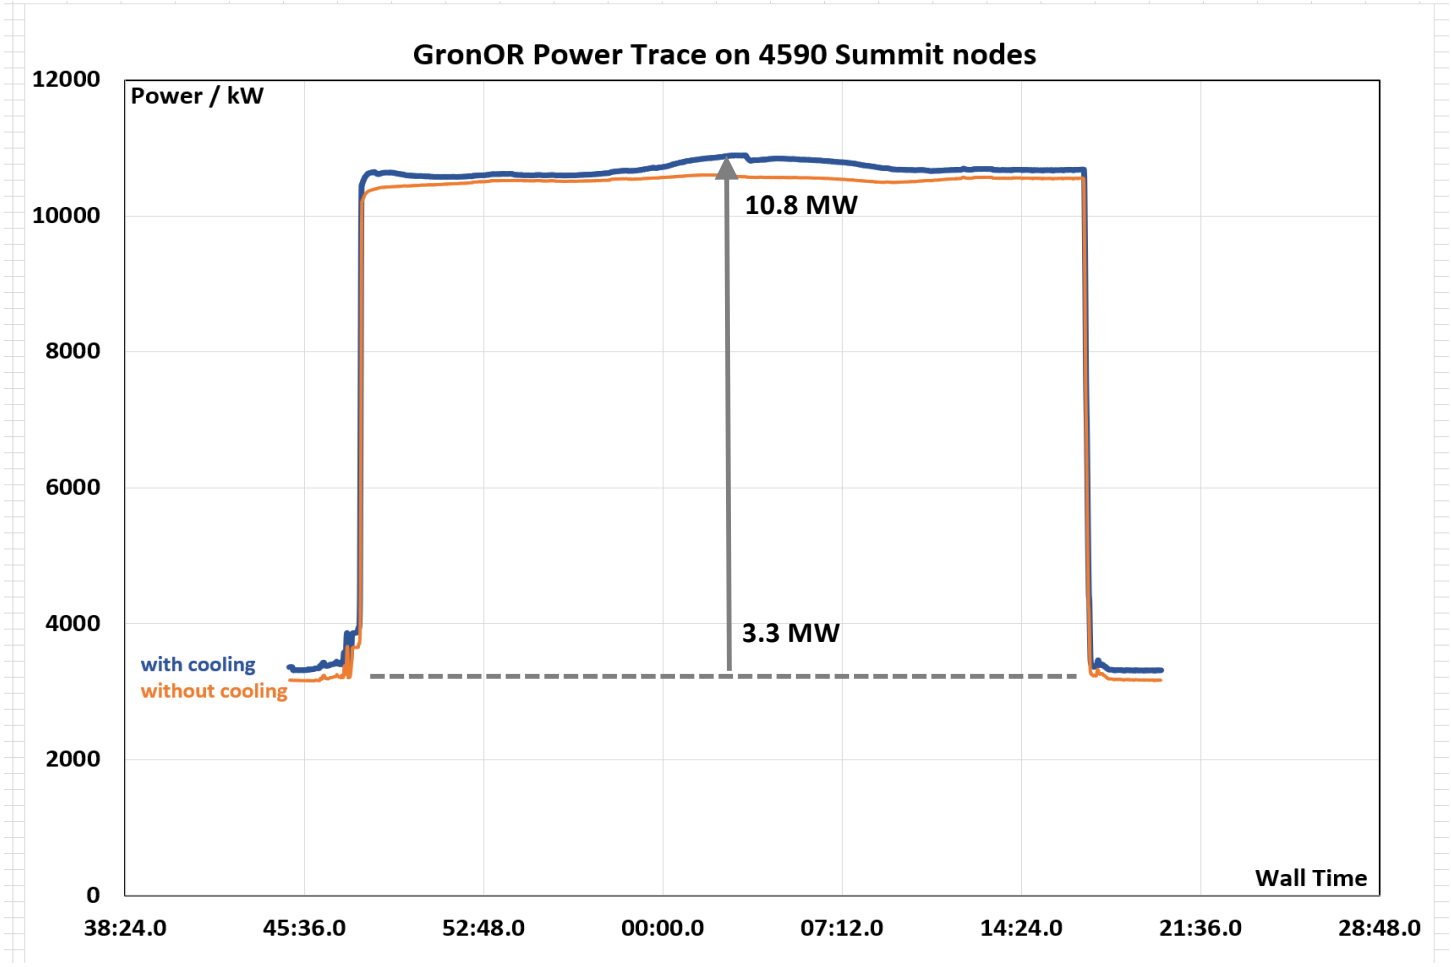

GronOR Power Trace on 4604 Summit nodes

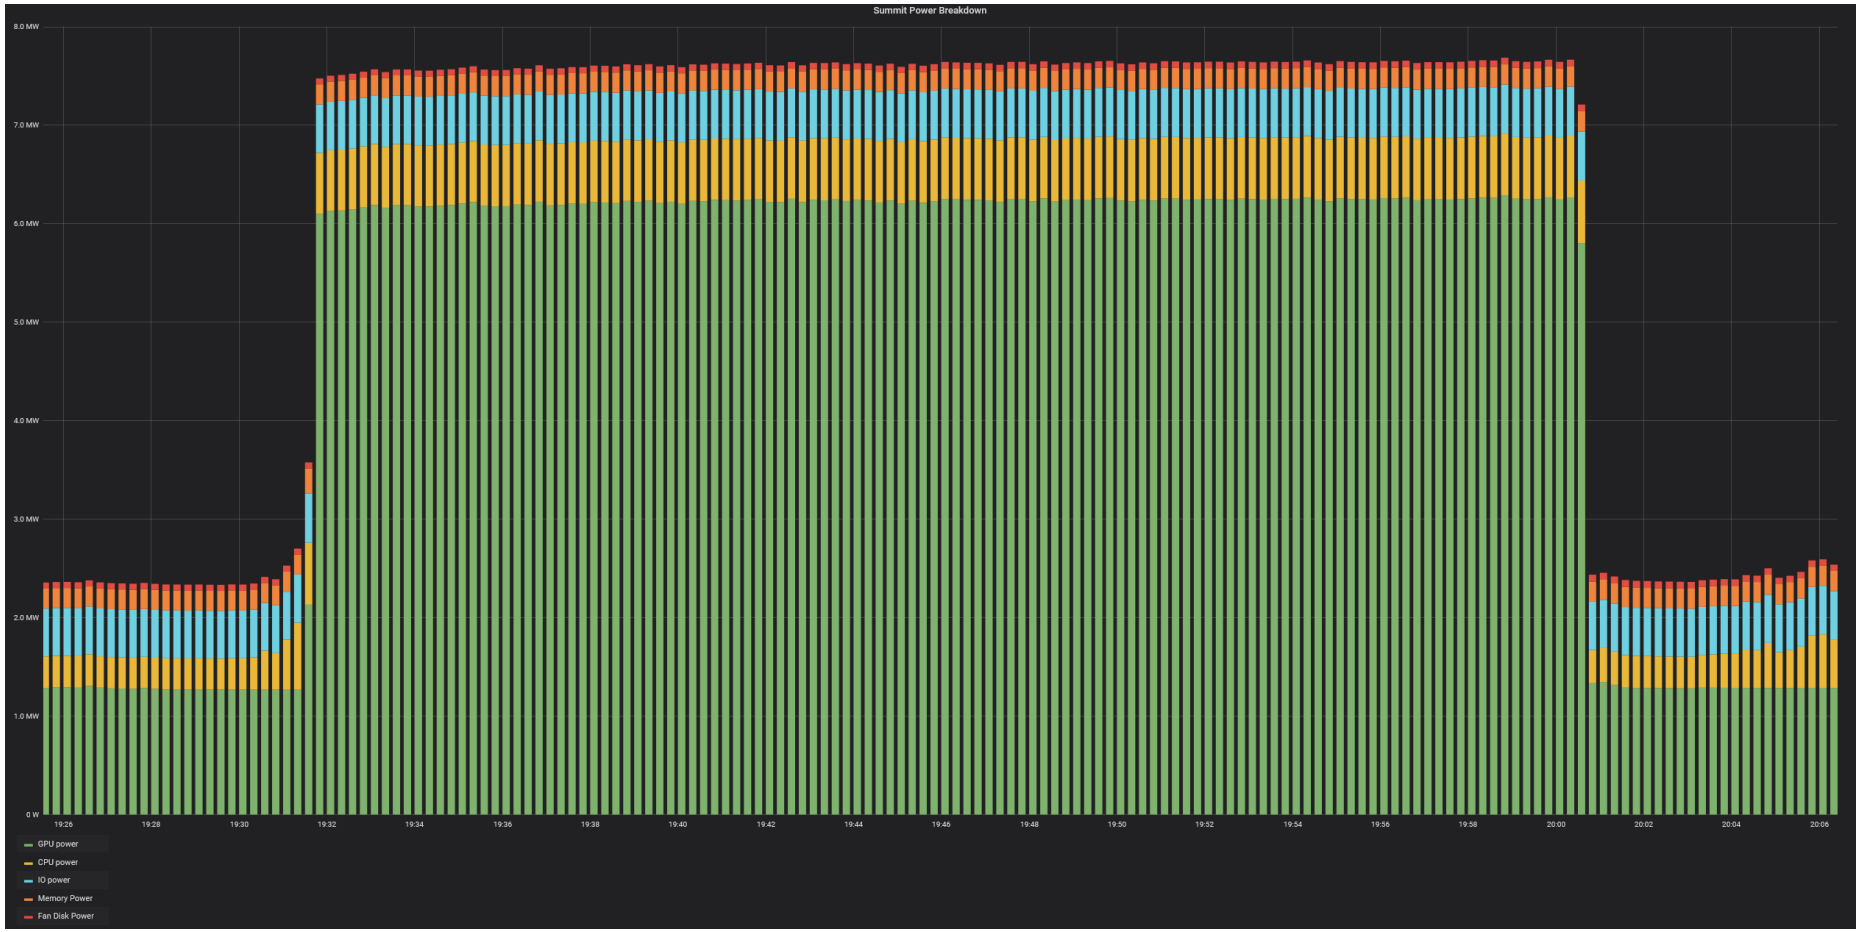

Grafana Dashboard on Summit

Indolonaphtyridine dimer benchmarks at CAS(4,4;4,4) on Summit

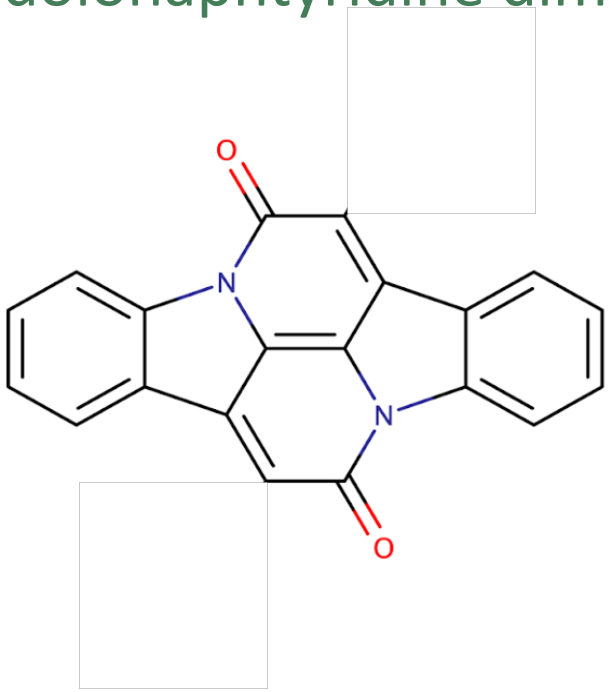

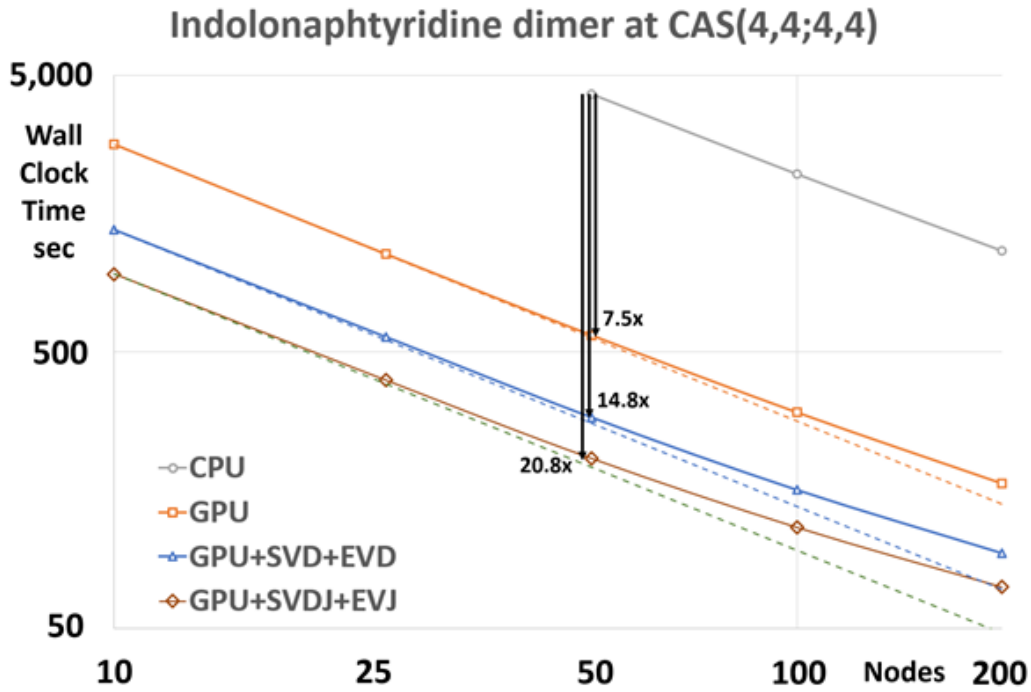

CPU: 30 ranks per node running on CPUs only;
GPU: 6 ranks per node using GPUs with SVD and EVD running on CPUs;
GPU+SVD+EVD: 6 ranks per node running on GPUs with QR-based CUSOLVER solvers GESVD and SYEVD;
GPU+SVDJ+EVDJ: 6 ranks per node running on GPUs with Jacobi iterative CUSOLVER solvers GESVDJ and SYEVJ;
Dashed lines indicate ideal linear scaling.

Indolonaphtyridine and Perylenediimide dimer benchmarks

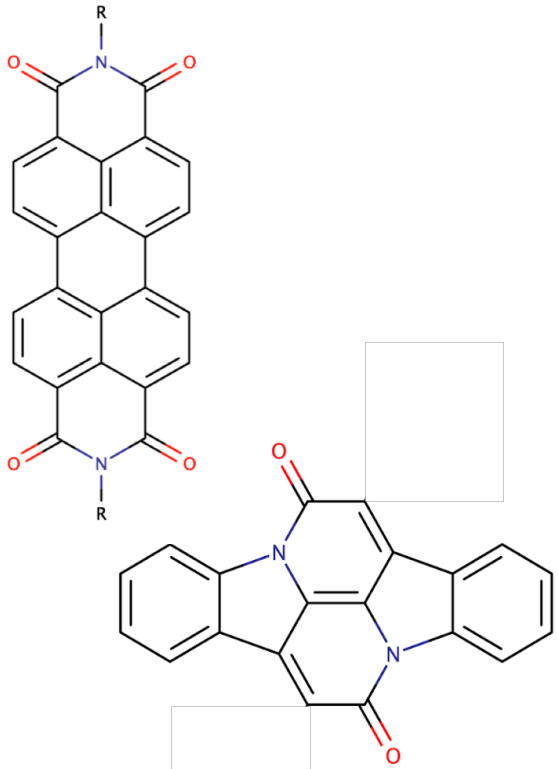

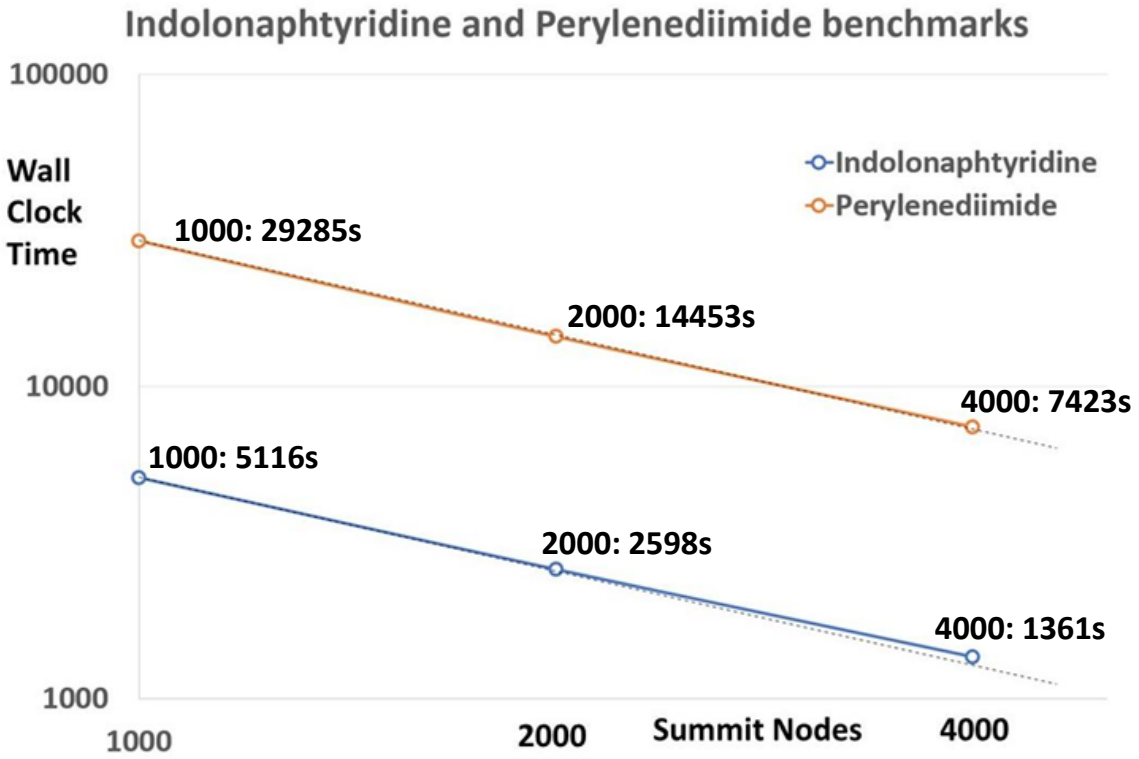

Benchmark timings for a perylenediimide and an indolonaphtyridine dimer at CAS(8,8;8,8) as a function of 1,000 to 4,000 Summit nodes used. The dashed lines indicate ideal linear scaling.

GronOR Acknowledgments

This work is part of the (Shell-NWO) research program of the Foundation for Fundamental Research on Matter (FOM), which is part of the Netherlands Organization for Scientific Research (NWO), and part of a European Joint Doctorate (EJD) in Theoretical Chemistry and Computational Modelling (TCCM), which has been financed under the framework of the Innovative Training Networks (ITN) of the MARIE Skłodowska-CURIE Actions (ITN-EJD-642294-TCCM)

This work used resources of the Oak Ridge Leadership Computing Facility (OLCF) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. The OLCF provided the opportunity to get early access to Summit through the Director's Discretionary program.

http://www.gronor.org

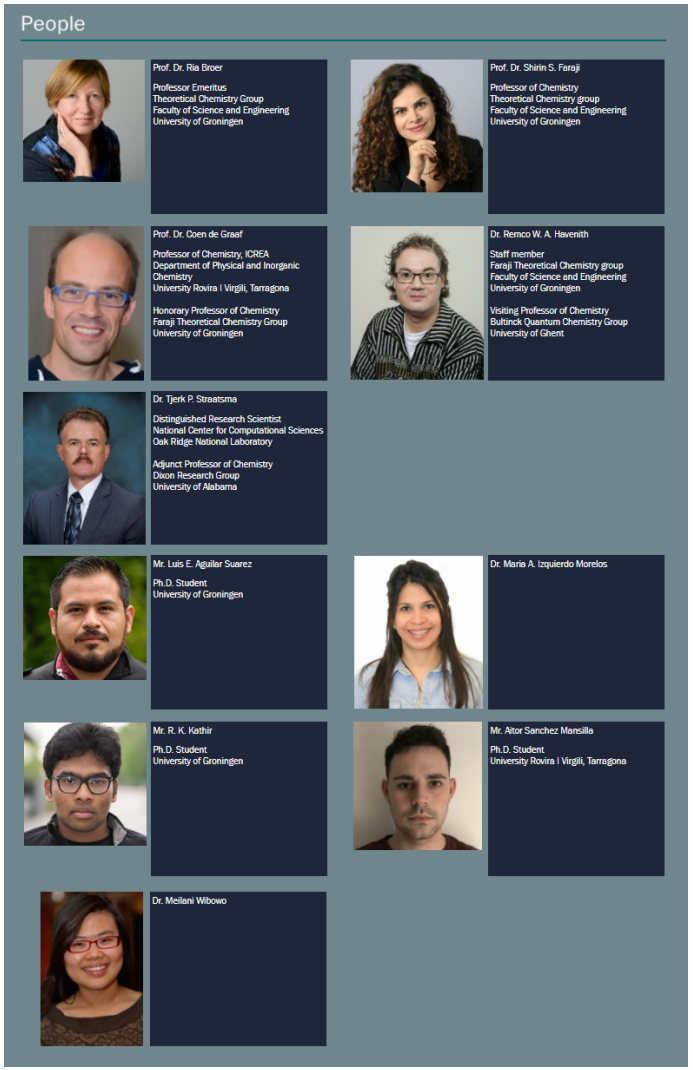

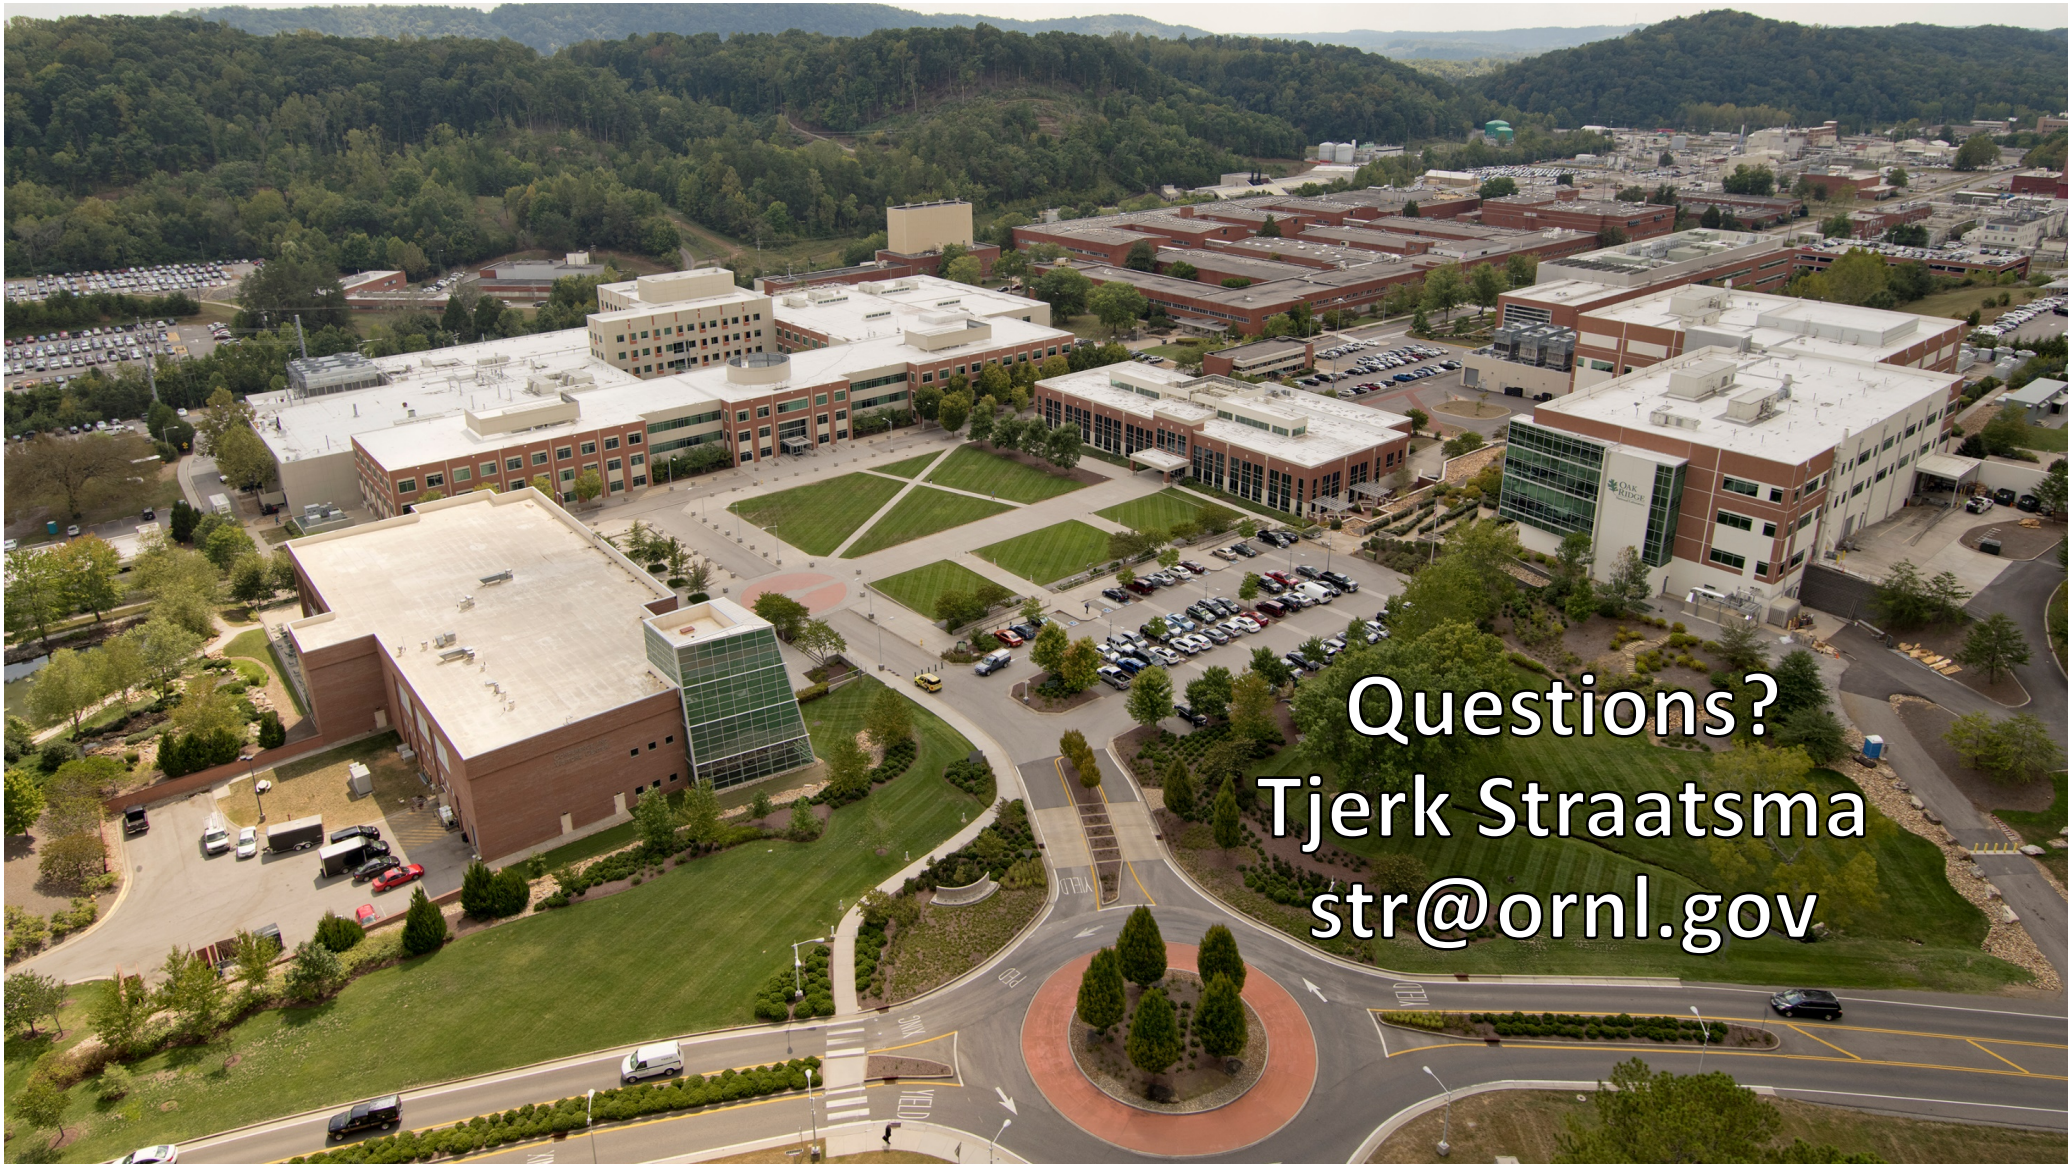
Questions?
Tjerk Straatsma
str@ornl.gov